

Approximate Clustering via Metric Partitioning*

Sayan Bandyapadhyay[†] and Kasturi Varadarajan[‡]

Department of Computer Science
University of Iowa, Iowa City, USA

Abstract

In this paper we consider two metric covering/clustering problems - *Minimum Cost Covering Problem* (MCC) and k -clustering. In the MCC problem, we are given two point sets X (clients) and Y (servers), and a metric on $X \cup Y$. We would like to cover the clients by balls centered at the servers. The objective function to minimize is the sum of the α -th power of the radii of the balls. Here $\alpha \geq 1$ is a parameter of the problem (but not of a problem instance). MCC is closely related to the k -clustering problem. The main difference between k -clustering and MCC is that in k -clustering one needs to select k balls to cover the clients.

For any $\varepsilon > 0$, we describe quasi-polynomial time $(1 + \varepsilon)$ approximation algorithms for both of the problems. However, in case of k -clustering the algorithm uses $(1 + \varepsilon)k$ balls. Prior to our work, a 3^α and a c^α approximation were achieved by polynomial-time algorithms for MCC and k -clustering, respectively, where $c > 1$ is an absolute constant. These two problems are thus interesting examples of metric covering/clustering problems that admit $(1 + \varepsilon)$ -approximation (using $(1 + \varepsilon)k$ balls in case of k -clustering), if one is willing to settle for quasi-polynomial time. In contrast, for the variant of MCC where α is part of the input, we show under standard assumptions that no polynomial time algorithm can achieve an approximation factor better than $O(\log |X|)$ for $\alpha \geq \log |X|$.

1998 ACM Subject Classification I.3.5 Computational Geometry and Object Modeling

Keywords and phrases Approximation Algorithms, Clustering, Covering, Probabilistic Partitions

Digital Object Identifier 10.4230/LIPIcs...

1 Introduction

We consider two metric covering/clustering problems. In the first problem, we are given two point sets X (clients) and Y (servers), and a metric d on $X \cup Y$. For $z \in X \cup Y$ and $r \geq 0$, the ball $B(z, r)$ centered at z and having radius $r \geq 0$ is the set $\{y \in X \cup Y \mid d(z, y) \leq r\}$. A *cover* for a subset $P \subseteq X$ is a set of balls, each centered at a point of Y , whose union contains P . The *cost* of a set $\mathcal{B} = \{B_1, \dots, B_k\}$ of balls, denoted by $\text{cost}(\mathcal{B})$, is $\sum_{i=1}^k r(B_i)^\alpha$, where $r(B_i)$ is the radius of B_i , and $\alpha \geq 1$ is a parameter of the problem (but not of a problem instance). The goal is to compute a minimum cost *cover* for the clients X . We refer to this problem as the *Minimum Cost Covering Problem* (MCC).

In the second problem, we are given a set X of n points, a metric d on X , and a positive integer k . Unlike in the case of MCC, here each ball is centered at a point in X .¹ The *cost*

* This material is based upon work supported by the National Science Foundation under Grant CCF-1318996

[†] sayan-bandyapadhyay@uiowa.edu

[‡] kasturi-varadarajan@uiowa.edu

¹ Our results do generalize to the problem where we distinguish between clients and servers as in the MCC.



© Sayan Bandyapadhyay and Kasturi Varadarajan;
licensed under Creative Commons License CC-BY

Leibniz International Proceedings in Informatics

LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

$\text{cost}(\mathcal{B})$ of a set \mathcal{B} of balls is defined exactly in the same way as in the case of MCC. The goal is to find a set \mathcal{B} of k balls whose union contains all the points in X and $\text{cost}(\mathcal{B})$ is minimized. We refer to this problem as k -clustering.

Inspired by applications in wireless networks, MCC has been well studied [21]. One can consider the points in Y as the potential locations of mobile towers and the points in X as the locations of customers. A tower can be configured in a way so that it can serve the customers lying within a certain distance. But the service cost increases with the distance served. The goal is to serve all the customers minimizing the total cost. For modelling the energy needed for wireless transmission, it is common to consider the value of α to be at least 1.

For the MCC problem with $\alpha = 1$, a primal-dual algorithm of Charikar and Panigrahy [9] leads to an approximation guarantee of 3; their result generalizes to $\alpha \geq 1$, with an approximation guarantee of 3^α . The problem is known to be NP-hard for $\alpha > 1$, even when X and Y are points in the Euclidean plane [2]. The case $\alpha = 1$ has received particular attention. The first PTAS for the Euclidean plane was designed by Lev-Tov and Peleg [21]. Later, Gibson *et. al* [16] have designed a polynomial time exact algorithm for this problem when X and Y are points in the plane, and the underlying distance function d is either the l_1 or l_∞ metric. For the l_2 metric they also get an exact algorithm if one assumes two candidate solutions can be compared efficiently; without this assumption, they get a $(1 + \varepsilon)$ approximation. Their algorithm is based on a separator theorem that, for any optimal solution, proves the existence of a balanced separator that intersects with at most 12 balls in the solution. In a different work they have also extended the exact algorithm to arbitrary metric spaces [15]. The running time is quasi-polynomial if the aspect ratio of the metric (ratio of maximum to minimum interpoint distance) is bounded by a polynomial in the number of points. When the aspect ratio is not polynomially bounded, they obtain a $(1 + \varepsilon)$ approximation in quasi-polynomial time. Their algorithms are based on a partitioning of the metric space that intersects a small number of balls in the optimal cover.

When $\alpha > 1$, the structure that holds for $\alpha = 1$ breaks down. It is no longer the case, even in the Euclidean plane, that there is a good separator (or partition) that intersects a small number of balls in an optimal solution. In the case $\alpha = 2$ and the Euclidean plane, the objective function models the total area of the served region, which arises in many practical applications. Hence this particular version has been studied in a series of works. Chuzhoy developed an unpublished 9-factor approximation algorithm for this version. Freund and Rawitz [14] present this algorithm and give a primal fitting interpretation of the approximation factor. Bilo *et. al* [8] have extended the techniques of Lev-Tov and Peleg [21] to get a PTAS that works for any $\alpha \geq 1$ and for any fixed dimensional Euclidean space. The PTAS is based on a sophisticated use of the *shifting strategy* which is a popular technique in computational geometry for solving problems in \mathbb{R}^d [12, 18]. For general metrics, however, the best known approximation guarantee for $\alpha > 1$ remains the already mentioned 3^α [9].

The k -clustering problem has applications in many fields including Data Mining, Machine Learning and Image Processing. Over the years it has been studied extensively from both theoretical and practical perspectives [8, 9, 11, 15, 16, 22]. The problem can be seen as a variant of MCC where $Y = X$ and at most k balls can be chosen to cover the points in X . As one might think, the constraint on the number of balls that can be used in k -clustering makes it relatively harder than MCC. Thus all the hardness results for MCC also hold for k -clustering. For $\alpha = 1$, Charikar and Panigrahy [9] present a polynomial time algorithm with an approximation guarantee of about 3.504. Gibson *et. al* [15, 16] obtain the same results for k -clustering with $\alpha = 1$ as the ones described for MCC, both in \mathbb{R}^d and arbitrary metrics.

Recently, Salavatipour and Behsaz [7] have obtained a polynomial time exact algorithm for $\alpha = 1$ and metrics of unweighted graphs, if we assume that no singleton clusters are allowed. However, in case of $\alpha > 1$ the best known approximation factor (in polynomial time) for general metrics is c^α , for some absolute constant $c > 1$; this follows from the analysis of Charikar and Panigrahy [9], who explicitly study only the case $\alpha = 1$. In fact, no better polynomial time approximation is known even for the Euclidean plane. We note that though the polynomial time algorithm in [8] yields a $(1 + \varepsilon)$ approximation for k -clustering in any fixed dimensional Euclidean space and for $\alpha \geq 1$, it can use $(1 + \varepsilon)k$ balls.

In addition to k -clustering many other clustering problems (k -means, k -center, k -median etc.) have been well studied [4, 10, 23, 17].

In this paper we address the following interesting question. Can the techniques employed by [8] for fixed dimensional Euclidean spaces be generalized to give $(1 + \varepsilon)$ approximation for MCC and k -clustering in any metric space? Our motivation for studying the problems in a metric context is partly that it includes two geometric contexts: (a) high dimensional Euclidean spaces; and (b) shortest path distance metric in the presence of polyhedral obstacles in \mathbb{R}^2 or \mathbb{R}^3 .

1.1 Our Results and Techniques

In this paper we consider the metric MCC and k -clustering with $\alpha \geq 1$. For any $\epsilon > 0$, we design a $(1 + \epsilon)$ -factor approximation algorithm for MCC that runs in quasi-polynomial time, that is, in $2^{(\log mn/\epsilon)^c}$ time, where $c > 0$ is a constant, $m = |Y|$, and $n = |X|$. We also have designed a similar algorithm for k -clustering that uses at most $(1 + \epsilon)k$ balls and yields a solution whose cost is at most $(1 + \epsilon)$ times the cost of an optimal k -clustering solution. The time complexity of the latter algorithm is also quasi-polynomial. As already noted, somewhat stronger guarantees are already known for the case $\alpha = 1$ of these problems [15], but the structural properties that hold for $\alpha = 1$ make it rather special.

The results in this paper should be compared with the polynomial time algorithms [9] that guarantee 3^α approximation for MCC and c^α approximation for k -clustering. The MCC and k -clustering are thus interesting examples of metric covering/clustering problems that admit $(1 + \epsilon)$ -approximation (using $(1 + \epsilon)k$ balls in case of k -clustering), if one is willing to settle for quasi-polynomial time. From this perspective, our results are surprising, as most of the problems in general metrics are APX-hard. The MCC and k -clustering are also examples where the techniques used in fixed dimensional Euclidean spaces generalize nicely to metric spaces. This is in contrast to the *facility location problem* [3].

The algorithms that we have designed for both of the problems use similar techniques that exploit the following key property of optimal covers: there are only a “small” number of balls whose radius is “large”. We can therefore afford to guess these balls by an explicit enumeration. However, there can be a “large” number of balls with “small” radius. To help ‘find’ these, we partition the metric space into blocks (or subsets) with at most half the original diameter, and recurse on each block. We have to pay a price for this recursion in the approximation guarantee. This price depends on the number of blocks in the partition that a small radius ball can intersect. (This is not an issue in the case $\alpha = 1$, where each ball that is not guessed intersects precisely one of the blocks [15].)

We are led to the following problem: is there a way to probabilistically partition a metric space into blocks of at most half the diameter, so that for any ball with “small” radius, the expected number of blocks that intersect the ball can be nicely bounded? The celebrated partitioning algorithms of Bartal [5] and Fakcharoenphol, Rao, and Talwar [13] guarantee that the probability that such a ball is intersected by two or more blocks is nicely bounded.

However, their bounds on the probability that a small ball is intersected do not directly imply a good bound on the expected number of blocks intersected by a small ball. Indeed, if one employs the partitioning algorithm of [13], the expected number of blocks intersected by a small ball can be quite “large”. Fortunately, the desired bound on the expectation can be shown to hold for the algorithm of Bartal [5], even though he did not study the expectation itself. We use a similar partitioning scheme and derive the expectation bound in Section 2, using an analysis that closely tracks previous work [1, 6, 19]. While the bound on the expectation is easily derived from previous work, our work is the first to study and fruitfully apply this bound.

The algorithms for MCC and k -clustering, which use the partitioning scheme of Section 2, are described in Section 3 and 4, respectively. In Section 5, we consider the approximability of a variant of the MCC where we allow α to be part of the input. For $\alpha \geq \log |X|$, we show, under standard complexity theoretic assumptions, that no polynomial (or quasi-polynomial) time algorithm for MCC can achieve an approximation factor better than $O(\log |X|)$. This partly explains the dependence on α of the running time of our algorithms.

2 The Partitioning Scheme

Let Z be a point set with an associated metric d , let $P \subseteq Z$ be a point set with at least 2 points, and $n \geq |P|$ be a parameter. For $Q \subseteq Z$, denote the maximum interpoint distance (or diameter) of Q by $\text{diam}(Q)$. Consider any partition of P into subsets (or blocks) $\{P_1, P_2, \dots, P_t\}$, where $2 \leq t \leq |P|$. Abusing notation, we will also view $\{P_1, P_2, \dots, P_t\}$ as a sequence of blocks. We say that P_i non-terminally (resp. terminally) intersects a ball B if P_i intersects B and it is not (resp. it is) the last set in the sequence P_1, P_2, \dots, P_t that intersects B . We would like to find a partition $\{P_1, P_2, \dots, P_t\}$ of P that ensures the following properties:

1. For each $1 \leq i \leq t$, $\text{diam}(P_i) \leq \text{diam}(P)/2$.
2. For any ball B (centered at some point in Z) of radius $r \leq \frac{\text{diam}(P)}{16 \log n}$, the expected size of the set $\{i | P_i \cap B \neq \emptyset\}$ is at most $1 + c \frac{r}{\text{diam}(P)} \log n$, where $c > 0$ is a constant. In other words, the expected number of blocks in the partition that intersect B is at most $1 + c \frac{r}{\text{diam}(P)} \log n$.
3. For any ball B (centered at some point in Z) of radius $r \leq \frac{\text{diam}(P)}{16 \log n}$, the expected number of blocks in the partition that non-terminally intersect B is at most $c \frac{r}{\text{diam}(P)} \log n$, where $c > 0$ is a constant.

We note that the second property follows from the third, as the number of blocks that intersect ball B is at most one more than the number of blocks that non-terminally intersect B . We will design a probabilistic partitioning algorithm that finds a partition with the desired properties.

2.0.1 The Partitioning Scheme of [13]

We first explain why the probabilistic partitioning algorithm of [13] does not achieve this guarantee. In this algorithm, we first pick a β uniformly at random from the interval $[\frac{\delta}{8}, \frac{\delta}{4}]$, where $\delta = \text{diam}(P)$. Also, let $\pi_1, \pi_2, \dots, \pi_p$ be a permutation of P chosen uniformly at random. We compute P_1, P_2, \dots, P_p in order as follows. Suppose we already computed P_1, \dots, P_{i-1} . We let

$$P_i = \{x \in P \setminus (P_1 \cup P_2 \cup \dots \cup P_{i-1}) : d(x, \pi_i) \leq \beta\}.$$

We will refer to P_i as π_i 's cluster. We return the partition $\{P_i \mid P_i \neq \emptyset\}$.

Consider the following weighted tree. Let vertex u be connected to vertices u_1, u_2, \dots, u_b using edges of weight $\frac{\delta}{16 \log n}$. Here, n, δ , and b are parameters. Let V_1, V_2, \dots, V_b be disjoint sets with b vertices each. For each i , u_i is connected to every vertex in V_i using edges of weight $\frac{\delta}{4} - \frac{\delta}{16 \log n}$. Finally, z is a new vertex that is connected to u using an edge of weight $\frac{3\delta}{4}$. Consider the metric induced by this weighted graph, and let P denote the vertex set. That is $P = \{u\} \cup \bigcup_i \{u_i\} \cup \bigcup_i V_i \cup \{z\}$. Set $n := |P|$ and note that $b = \Theta(\sqrt{n})$. Also, $\delta = \text{diam}(P)$.

Let $B = B(u, r)$ where $r = \frac{\delta}{16 \log n}$. Notice that the ball B consists of the points $\{u, u_1, u_2, \dots, u_b\}$. Consider running the probabilistic partitioning algorithm of [13], described above, on P . We argue that the expected number of blocks in the output partition that intersect B is $\Omega(\frac{\sqrt{n}}{\log n})$, which is asymptotically larger than $1 + c \frac{r}{\text{diam}(P)} \log n = O(1)$.

Fix $1 \leq i \leq b$. We observe that in the output partition, u_i belongs to the cluster of some point in $\{u, u_1, u_2, \dots, u_b\}$ or of some point in V_i . Call u_i *isolated* if u_i belongs to the cluster of some point in V_i . If u_i is isolated, the cluster containing u_i does not contain any of the other u_j . Thus, the number of blocks of the output partition that intersect B is at least the number of vertices in $\{u_1, u_2, \dots, u_b\}$ that are isolated.

Note that u_i is isolated if the following two events occur: (a) $\beta \in [\frac{\delta}{4} - \frac{\delta}{16 \log n}, \frac{\delta}{4}]$; (b) some vertex in V_i appears in the random permutation $\pi_1, \pi_2, \dots, \pi_p$ before all vertices in $\{u, u_1, \dots, u_b\}$. The probability of these two events occurring is $\Omega(\frac{1}{\log n})$. It follows that the expected number of isolated vertices, and thus the expected number of blocks in the partition that intersect B , is $\Omega(\frac{b}{\log n}) = \Omega(\frac{\sqrt{n}}{\log n})$.

2.0.2 Probability Distribution

Before describing our partitioning algorithm, we consider a probability distribution that it uses. Given a positive real δ and an integer $k \geq 2$, the distribution is denoted by $\text{dist}(\delta, k)$. The probability density function (pdf) f of $\text{dist}(\delta, k)$ is the following:

$$f(x) = \begin{cases} 0 & \text{if } x < \delta/8 \text{ and } x > \delta/4 \\ \frac{8 \log k}{\delta} \frac{1}{2^i} & \text{if } \frac{\delta}{8} + (i-1) \frac{\delta}{8 \log k} \leq x < \frac{\delta}{8} + i \frac{\delta}{8 \log k} \text{ for } 1 \leq i \leq \log k - 1 \\ \frac{8 \log k}{\delta} \frac{2}{k} & \text{if } \frac{\delta}{4} - \frac{\delta}{8 \log k} \leq x \leq \frac{\delta}{4} \end{cases}$$

The following observation shows that f is indeed a density function.

► **Observation 2.1.** f is a probability density function.

Proof. It is sufficient to show that f satisfies the two properties of density function. As δ and k are nonnegative it is easy to see that $f(x) \geq 0$ for $x \in (-\infty, +\infty)$. Also,

$$\begin{aligned} \int_{-\infty}^{\infty} f(x) dx &= \int_{\frac{\delta}{8}}^{\frac{\delta}{8} + \frac{\delta}{8 \log k}} \frac{8 \log k}{\delta} \frac{1}{2} dx + \int_{\frac{\delta}{8} + \frac{\delta}{8 \log k}}^{\frac{\delta}{8} + \frac{2\delta}{8 \log k}} \frac{8 \log k}{\delta} \frac{1}{2^2} dx \\ &\quad + \dots + \int_{\frac{\delta}{8} + \frac{(\log k - 2)\delta}{8 \log k}}^{\frac{\delta}{8} + \frac{(\log k - 1)\delta}{8 \log k}} \frac{8 \log k}{\delta} \frac{1}{2^{\log k - 1}} dx + \int_{\frac{\delta}{4} - \frac{\delta}{8 \log k}}^{\frac{\delta}{4}} \frac{8 \log k}{\delta} \frac{2}{k} dx \\ &= \frac{1}{2} \sum_{i=0}^{\log k - 2} \frac{1}{2^i} + \frac{2}{k} = 1 - \left(\frac{1}{2}\right)^{\log k - 1} + \frac{2}{k} = 1 \end{aligned}$$

◀

Consider the interval $[\frac{\delta}{8}, \frac{\delta}{4}]$. Now divide the interval into $\log k$ subintervals of equal length. The i^{th} interval for $1 \leq i \leq \log k - 1$ is defined as $[\frac{\delta}{8} + (i-1)\frac{\delta}{8\log k}, \frac{\delta}{8} + i\frac{\delta}{8\log k})$. The last interval is $[\frac{\delta}{4} - \frac{\delta}{8\log k}, \frac{\delta}{4}]$. Denote the j^{th} interval by I_j for $1 \leq j \leq \log k$.

To sample a β according to $\text{dist}(\delta, k)$, we first pick one of these intervals from the distribution that assigns a probability of $1/2^j$ to I_j for $1 \leq j \leq \log k - 1$, and a probability of $2/k$ to $I_{\log k}$. Having picked an interval I_j , we generate β uniformly at random from it.

Now we discuss a vital property of the distribution $\text{dist}(\delta, k)$ which we use in the analysis of the partitioning algorithm. For an event E , let $\Pr[E]$ denotes the probability that E occurs. Now consider the following random process. We sample a value p from $\text{dist}(\delta, k)$. Let E_j denotes the event that $p \in I_j$. Then we have the following observation.

► **Observation 2.2.** $\Pr[E_j] = \sum_{i=j+1}^{\log k} \Pr[E_i]$ for $1 \leq j \leq \log k - 1$.

Proof. The proof follows from the definition of the pdf of $\text{dist}(\delta, k)$.

$$\begin{aligned} \Pr[E_j] &= \int_{\frac{\delta}{8} + (j-1)\frac{\delta}{8\log k}}^{\frac{\delta}{8} + j\frac{\delta}{8\log k}} \frac{8\log k}{\delta} \frac{1}{2^j} dx = \frac{1}{2^j} = \frac{1}{2^j} - \frac{2}{k} + \frac{2}{k} = \frac{1}{2^j} \left(1 - \frac{2^{j+1}}{k}\right) + \frac{2}{k} \\ &= \frac{1}{2^j} \left(1 - \frac{1}{2^{\log k - j - 1}}\right) + \frac{2}{k} = \frac{1}{2^{j+1}} \frac{(1 - \frac{1}{2^{\log k - j - 1}})}{\frac{1}{2}} + \frac{2}{k} \\ &= \frac{1}{2^{j+1}} \left(\frac{1}{2^0} + \frac{1}{2^1} + \dots + \frac{1}{2^{\log k - j - 2}}\right) + \frac{2}{k} \\ &= \frac{1}{2^{j+1}} + \dots + \frac{1}{2^{\log k - 1}} + \frac{2}{k} \\ &= \Pr[E_{j+1}] + \dots + \Pr[E_{\log k - 1}] + \Pr[E_{\log k}] = \sum_{i=j+1}^{\log k} \Pr[E_i] \end{aligned}$$

◀

2.0.3 Partitioning Algorithm

Now we move on to the partitioning algorithm, which we recall, is given Z , a metric d on Z , $P \subseteq Z$, and a parameter $n \geq |P|$. The procedure $\text{RAND-PARTITION}(P)$ described below (as Algorithm 1) takes a point set $P \subseteq Z$ as input and outputs a partition with at most $|P|$ subsets. Suppose that $P = \{p_1, \dots, p_{|P|}\}$. The algorithm then generates $P_1, P_2, \dots, P_{|P|}$ in order via the for loop. Suppose that P_1, P_2, \dots, P_{i-1} have already been constructed, and $Q = P \setminus (P_1 \cup P_2 \cup \dots \cup P_{i-1})$. To construct P_i , the procedure samples a β_i from $\text{dist}(\text{diam}(P), n)$. The choice of β_i is done independently of the choices of the other points. Then P_i is set to $\{x \in Q \mid d(x, p_i) \leq \beta_i\}$. Note that this is done in the i^{th} iteration of the for loop. (Note that p_i might not be assigned to P_i , as it could already be assigned to some other subset.)

We show that $\text{RAND-PARTITION}(P)$ satisfies the two guarantees mentioned before. To see that the first guarantee, note that the β values are chosen from the distribution $\text{dist}(\text{diam}(P), n)$ which ensures that $\beta_i \leq \text{diam}(P)/4$ for $1 \leq i \leq p$. Now each point in a subset P_i is at a distance at most β_i from p_i . Thus by triangle inequality $\text{diam}(P_i) \leq \text{diam}(P)/2$. In the next lemma we will show that the second guarantee also holds. Before that we have a definition.

We say that P_i non-terminally (resp. terminally) intersects $B(y, r)$ if P_i intersects $B(y, r)$ and it is not (resp. it is) the last set in the sequence P_1, P_2, \dots, P_p that intersects $B(y, r)$.

Algorithm 1 RAND-PARTITION(P)**Require:** A subset $P = \{p_1, \dots, p_{|P|}\} \subseteq Z$ **Ensure:** A partition of P

- 1: $p \leftarrow |P|$
- 2: $Q \leftarrow P$
- 3: **for** $i = 1$ to p **do**
- 4: sample a β_i from $\text{dist}(\text{diam}(P), n)$ corresponding to p_i
- 5: $P_i \leftarrow \{x \in Q \mid d(x, p_i) \leq \beta_i\}$
- 6: $Q \leftarrow Q \setminus P_i$
- 7: **return** $\{P_i \mid P_i \neq \emptyset \text{ and } 1 \leq i \leq p\}$

► **Lemma 1.** (*Partitioning Lemma*) *There is a constant c such that for any ball $B(y, r)$ with $r \leq \frac{\text{diam}(P)}{16 \log n}$, the expected number of blocks in the output of RAND-PARTITION(P) that intersect $B(y, r)$ is at most $1 + c \frac{r}{\text{diam}(P)} \log n$. Moreover, the expected number of blocks that intersect the ball non-terminally is at most $c \frac{r}{\text{diam}(P)} \log n$.*

The intuition for the lemma is as follows. Consider the beginning of the i 'th iteration of the for loop and assume that ball $B(y, r)$ is not fully contained in the union of the previously constructed blocks P_1, \dots, P_{i-1} . Then, considering the choice of β_i , the probability that $B(p_i, \beta_i)$ fully contains the ball $B(y, r)$ is nearly as large as the probability that $B(p_i, \beta_i)$ intersects $B(y, r)$. If $B(p_i, \beta_i)$ fully contains the ball $B(y, r)$, then of course none of the blocks $P_{i+1}, P_{i+2}, \dots, P_p$ intersect $B(y, r)$. We now proceed to the proof.

Proof. For a point $x \in Z$ and subset $Q \subseteq Z$, let $d_{\min}(x, Q) = \min_{q \in Q} d(x, q)$ and $d_{\max}(x, Q) = \max_{q \in Q} d(x, q)$. Fix the ball $B(y, r)$ with $r \leq \frac{\text{diam}(P)}{16 \log n}$. For each $1 \leq i \leq p$, consider the indicator random variable T_i defined as follows:

$$T_i = \begin{cases} 1 & \text{if } P_i \text{ intersects } B(y, r) \\ 0 & \text{otherwise} \end{cases}$$

Let the random variable $T = \sum_{i=1}^p T_i$ be the number of subsets that the ball intersects. Then $E[T] = \sum_{i=1}^p E[T_i] = \sum_{i=1}^p \Pr[P_i \text{ intersects } B(y, r)]$.

Clearly, there is at most one P_i that is the last one that intersects $B(y, r)$. Thus,

$$\sum_{i=1}^p \Pr[P_i \text{ intersects } B(y, r)] \leq 1 + \sum_{i=1}^p \Pr[P_i \text{ non-terminally intersects } B(y, r)].$$

Let $x_i = d_{\min}(p_i, B(y, r))$ and $y_i = d_{\max}(p_i, B(y, r))$. By the triangle inequality, $y_i - x_i \leq 2r$. Denote by (S^i) the event that β_i lands in the interval $[x_i, y_i]$. Note that for P_i to non-terminally intersect $B(y, r)$, the event (S^i) must occur. Thus, if the interval $[x_i, y_i]$ does not intersect the interval $[\frac{\text{diam}(P)}{8}, \frac{\text{diam}(P)}{4}]$, then $\Pr[P_i \text{ non-terminally intersects } B(y, r)] = 0$.

We therefore turn to the case where $[x_i, y_i]$ does intersect the interval $[\frac{\text{diam}(P)}{8}, \frac{\text{diam}(P)}{4}]$. Recall that in defining the probability distribution $\text{dist}(\text{diam}(P), n)$, we have divided the latter interval into $\log n$ subintervals $I_1, I_2, \dots, I_{\log n}$ of equal length. Denote by a_l the probability

$$\Pr[\text{a random sample drawn from } \text{dist}(\text{diam}(P), n) \text{ belongs to } I_l].$$

For convenience, define $I_{\log n+1} = [\text{diam}(P)/4, \infty)$ and $a_{\log n+1} = 0$.

XX:8 Approximate Clustering via Metric Partitioning

Let I_{l_i} be the subinterval that contains x_i . (In case $x_i < \frac{\text{diam}(P)}{8}$, let $l_i = 1$.) The length of $[x_i, y_i]$ is at most $2r$, $2r \leq \text{diam}(P)/8 \log n$, and the length of each of the subintervals is $\text{diam}(P)/8 \log n$. Thus $[x_i, y_i]$ can intersect at most one more subinterval, and this is I_{l_i+1} . Let r_1 and r_2 be the length of $I_{l_i} \cap [x_i, y_i]$ and $I_{l_i+1} \cap [x_i, y_i]$ respectively. Note that $r_1 + r_2 \leq y_i - x_i \leq 2r$.

To bound $\Pr[P_i \text{ non-terminally intersects } B(y, r)]$, we now have two cases. We say that p_i is *far* (from the ball $B(y, r)$) if $l_i \in \{\log n - 1, \log n\}$. We say that p_i is *near* if $1 \leq l_i \leq \log n - 2$.

Case 1: p_i is far. In this case $a_{l_i}, a_{l_i+1} \leq \frac{2}{n}$. Thus

$$\begin{aligned} \Pr[S^i] &\leq \Pr[\beta_i \text{ lands in } I_{l_i} \cap [x_i, y_i]] + \Pr[\beta_i \text{ lands in } I_{l_i+1} \cap [x_i, y_i]] \\ &\leq \frac{r_1}{\text{diam}(P)/8 \log n} a_{l_i} + \frac{r_2}{\text{diam}(P)/8 \log n} a_{l_i+1} \\ &\leq \frac{2r}{\text{diam}(P)/8 \log n} \cdot \frac{2}{n} \\ &= \frac{32r \log n}{n \cdot \text{diam}(P)}. \end{aligned}$$

Thus, $\Pr[P_i \text{ non-terminally intersects } B(y, r)] \leq \Pr[S^i] \leq \frac{32r \log n}{n \cdot \text{diam}(P)}$.

Case 2: p_i is near. For such a p_i we have the following crucial observation.

► **Claim 2.1.** $\Pr[P_i \text{ non-terminally intersects } B(y, r)] \leq \frac{32r \log n}{\text{diam}(P)} \Pr[P_i \text{ terminally intersects } B(y, r)]$.

Proof. Suppose that P_1, P_2, \dots, P_{i-1} have been chosen and $B(y, r) \subseteq P_1 \cup P_2 \cup \dots \cup P_{i-1}$. Conditioned on such a history, we have $\Pr[P_i \text{ non-terminally intersects } B(y, r)] = \Pr[P_i \text{ terminally intersects } B(y, r)] = 0$ and the claimed inequality holds.

Now suppose that $B(y, r) \setminus (P_1 \cup P_2 \cup \dots \cup P_{i-1}) \neq \emptyset$. Let us condition on such a history. Then, P_i terminally intersects $B(y, r)$ if β_i lands in $I_{l_i+2} \cup I_{l_i+3} \cup \dots \cup I_{\log n}$. Thus, using Observation 2.2,

$$\Pr[P_i \text{ terminally intersects } B(y, r)] \geq a_{l_i+2} + a_{l_i+3} + \dots + a_{\log n} = a_{l_i+1}.$$

On the other hand,

$$\begin{aligned} \Pr[P_i \text{ non-terminally intersects } B(y, r)] &\leq \Pr[S^i \mid \beta_i \in I_{l_i} \cup I_{l_i+1}] \cdot \Pr[\beta_i \in I_{l_i} \cup I_{l_i+1}] \\ &\leq \left(\frac{2}{3} \frac{r_1}{\text{diam}(P)/8 \log n} + \frac{1}{3} \frac{r_2}{\text{diam}(P)/8 \log n} \right) \cdot (a_{l_i} + a_{l_i+1}) \\ &\leq \frac{32 \cdot r \log n}{3 \cdot \text{diam}(P)} \cdot 3a_{l_i+1} \\ &\leq \frac{32 \cdot r \log n}{\text{diam}(P)} \cdot \Pr[P_i \text{ terminally intersects } B(y, r)]. \end{aligned}$$

◀

Hence the expected number of subsets that intersect $B(y, r)$ non-terminally is

$$\begin{aligned}
&\leq \sum_{i=1}^p Pr[P_i \text{ non-terminally intersects } B(y, r)] \\
&\leq \sum_{i: p_i \text{ is far}} Pr[P_i \text{ non-terminally intersects } B(y, r)] \\
&\quad + \sum_{i: p_i \text{ is near}} Pr[P_i \text{ non-terminally intersects } B(y, r)] \\
&\leq \frac{32r \log n}{\text{diam}(P)} \sum_{i=1}^p \frac{1}{n} + \frac{32r \log n}{\text{diam}(P)} \sum_{i=1}^p Pr[P_i \text{ terminally intersects } B(y, r)] \\
&\leq c \frac{r}{\text{diam}(P)} \log n.
\end{aligned}$$

For the last inequality, we used the fact that $\sum_{i=1}^p Pr[P_i \text{ terminally intersects } B(y, r)] = 1$, since there is exactly one P_i that terminally intersects $B(y, r)$. Putting the two cases together, we have

$$E[T] \leq 1 + c \frac{r}{\text{diam}(P)} \log n.$$

◀

We conclude by summarizing the result.

► **Theorem 2.** *Let Z be a point set with an associated metric d , let $P \subseteq Z$ be a point set with at least 2 points, and $n \geq |P|$ be a parameter. There is a polynomial-time probabilistic algorithm $\text{RAND-PARTITION}(P)$ that partitions P into blocks $\{P_1, P_2, \dots, P_t\}$ and has the following guarantees:*

1. *For each $1 \leq i \leq t$, $\text{diam}(P_i) \leq \text{diam}(P)/2$.*
2. *There is a constant $c > 0$ so that for any ball B (centered at some point in Z) of radius $r \leq \frac{\text{diam}(P)}{16 \log n}$, the expected size of the set $\{i | P_i \cap B \neq \emptyset\}$ is at most $1 + c \frac{r}{\text{diam}(P)} \log n$ and the expected number of blocks that non-terminally intersect B is at most $c \frac{r}{\text{diam}(P)} \log n$.*

3 Algorithm for MCC

We now describe our $(1 + \epsilon)$ -factor approximation algorithm for the MCC problem. Recall that we are given a set X of clients, a set Y of servers, and a metric d on $X \cup Y$. We wish to compute a cover for X with minimum cost. Let $m = |Y|$ and $n = |X|$.

For $P \subseteq X$, let $\text{opt}(P)$ denote some optimal cover for P . Denote by $\text{cost}(B)$ the cost of a ball B (the α -th power of B 's radius) and by $\text{cost}(\mathcal{B})$ the cost $\sum_{B \in \mathcal{B}} \text{cost}(B)$ of a set \mathcal{B} of balls.

To compute a cover for P , our algorithm first guesses the set $\mathcal{Q} \subseteq \text{opt}(P)$ consisting of all the large balls in $\text{opt}(P)$. As we note in the structure lemma below, we may assume that the number of large balls in $\text{opt}(P)$ is small. We then use the algorithm of Theorem 2 to partition P into $\{P_1, P_2, \dots, P_t\}$. For each $1 \leq i \leq t$, we recursively compute a cover for the set $P'_i \subseteq P_i$ of points not covered by \mathcal{Q} .

To obtain an approximation guarantee for this algorithm, we use the guarantees of Theorem 2. With this overview, we proceed to the structure lemma and a complete description of the algorithm.

3.1 A Structure Lemma

It is not hard to show that for any $\gamma \geq 1$ and $P \subseteq X$ such that $\text{diam}(P)$ is at least a constant factor of $\text{diam}(X \cup Y)$, $\text{opt}(P)$ contains at most $(c/\gamma)^\alpha$ balls of radius at least $\text{diam}(P)/\gamma$. Here c is some absolute constant. The following structural lemma extends this fact.

► **Lemma 3.** *Let $P \subseteq X$, $0 < \lambda < 1$ and $\gamma \geq 1$, and suppose that $\text{opt}(P)$ does not contain any ball of radius greater than or equal to $2\alpha \cdot \text{diam}(P)/\lambda$. Then the number of balls in $\text{opt}(P)$ of radius greater than or equal to $\text{diam}(P)/\gamma$ is at most $c(\lambda, \gamma) := (9\alpha\gamma/\lambda)^\alpha$.*

Proof. Suppose that $\text{opt}(P)$ does not contain any ball of radius greater than or equal to $2\alpha \cdot \text{diam}(P)/\lambda$. Note that each ball in $\text{opt}(P)$ intersects P and has radius at most $2\alpha \cdot \text{diam}(P)/\lambda$. Thus the point set $\{z \in X \cup Y \mid z \in B \text{ for some } B \in \text{opt}(P)\}$ has diameter at most $\text{diam}(P) + 8\alpha \cdot \text{diam}(P)/\lambda \leq 9\alpha \cdot \text{diam}(P)/\lambda$. It follows that there is a ball centered at a point in Y , with radius at most $9\alpha \cdot \text{diam}(P)/\lambda$ that contains P .

Let t denote the number of balls in $\text{opt}(P)$ of radius greater than or equal to $\text{diam}(P)/\gamma$. By optimality of $\text{opt}(P)$, we have $t \cdot (\text{diam}(P)/\gamma)^\alpha \leq (9\alpha \cdot \text{diam}(P)/\lambda)^\alpha$. Thus $t \leq (9\alpha\gamma/\lambda)^\alpha$. ◀

3.2 The Algorithm

We may assume that the minimum distance between two points in X is 1. Let $L = 1 + \log(\text{diam}(X))$. As we want a $(1 + \varepsilon)$ -approximation, we fix a parameter $\lambda = \varepsilon/2L$. Let $\gamma = \frac{c \log n}{\lambda}$, where c is the constant in Theorem 2. Denote \mathcal{D} to be the set of balls such that each ball is centered at a point of $y \in Y$ and has radius $r = d(x, y)$ for some $x \in X$. We note that for any $P \subseteq X$, any ball in $\text{opt}(P)$ must belong to this set. Note that $|\mathcal{D}| \leq mn$. Recall that $c(\lambda, \gamma) = (9\alpha\gamma/\lambda)^\alpha$.

With this terminology, the procedure $\text{POINT-COVER}(P)$ described as Algorithm 2 returns a cover of $P \subseteq X$. If $|P|$ is smaller than some constant, then the procedure returns an optimal solution by searching all covers with a constant number of balls. In the general case, one candidate solution is the best single ball solution. For the other candidate solutions, the procedure first computes a partition $\{P_1, \dots, P_\tau\}$ of P , using the $\text{RAND-PARTITION}(P)$ procedure. Here $\text{RAND-PARTITION}(P)$ is called with $Z = X \cup Y$ and $n = |X| \geq |P|$. Then it iterates over all possible subsets of \mathcal{D} of size at most $c(\lambda, \gamma)$ containing balls of radius greater than $\text{diam}(P)/\gamma$. For each such subset \mathcal{Q} and $1 \leq i \leq \tau$, it computes the set $P'_i \subseteq P_i$ of points not covered by \mathcal{Q} . It then makes recursive calls and generates the candidate solution $\mathcal{Q} \cup \bigcup_{i=1}^\tau \text{POINT-COVER}(P'_i)$. Note that all the candidate solutions are actually valid covers for P . Among these candidate solutions the algorithm returns the best solution.

Our overall algorithm for MCC calls the procedure $\text{POINT-COVER}(X)$ to get a cover of X .

3.3 Approximation Guarantee

For $P \subseteq X$, let $\text{level}(P)$ denote the smallest non-negative integer i such that $\text{diam}(P) < 2^i$. As the minimum interpoint distance in X is 1, $\text{level}(P) = 0$ if and only if $|P| \leq 1$. Note that $\text{level}(X) \leq L$.

The following lemma bounds the quality of the approximation of our algorithm.

► **Lemma 4.** *$\text{POINT-COVER}(P)$ returns a solution whose expected cost is at most $(1 + \lambda)^l \text{cost}(\text{opt}(P))$, where $l = \text{level}(P)$.*

Algorithm 2 POINT-COVER(P)**Require:** A subset $P \subseteq X$.**Ensure:** A cover of the points in P .

```

1: if  $|P|$  is smaller than some constant  $\kappa$  then
2:   return a minimum solution by checking all covers with at most  $\kappa$  balls.
3:  $\text{sol} \leftarrow$  the best cover with one ball
4:  $\text{cost} \leftarrow \text{cost}(\text{sol})$ 
5: Let  $\{P_1, \dots, P_\tau\}$  be the set of nonempty subsets returned by RAND-PARTITION( $P$ )
6: Let  $\mathcal{B}$  be the set of balls in  $\mathcal{D}$  having radius greater than  $\frac{\text{diam}(P)}{\gamma}$ 
7: for each  $\mathcal{Q} \subseteq \mathcal{B}$  of size at most  $c(\lambda, \gamma)$  do
8:   for  $i = 1$  to  $\tau$  do
9:     Let  $P'_i = \{p \in P_i \mid p \notin \bigcup_{B \in \mathcal{Q}} B\}$ 
10:     $\mathcal{Q}' \leftarrow \mathcal{Q} \cup \bigcup_{i=1}^\tau \text{POINT-COVER}(P'_i)$ 
11:    if  $\text{cost}(\mathcal{Q}') < \text{cost}$  then
12:       $\text{cost} \leftarrow \text{cost}(\mathcal{Q}')$ 
13:       $\text{sol} \leftarrow \mathcal{Q}'$ 
14: return  $\text{sol}$ 

```

Proof. We prove this lemma using induction on l . If $l = 0$, then $|P| \leq 1$ and POINT-COVER(P) returns an optimal solution, whose cost is $\text{cost}(\text{opt}(P))$. Thus assume that $l \geq 1$ and the statement is true for subsets having level at most $l - 1$. Let $P \subseteq X$ be a point set with $\text{level}(P) = l$. If $|P|$ is smaller than the constant threshold κ , POINT-COVER(P) returns an optimal solution. So we may assume that $|P|$ is larger than this threshold. We have two cases.

Case 1: There is some ball in $\text{opt}(P)$ whose radius is at least $2\alpha \cdot \text{diam}(P)/\lambda$. Let B denote such a ball and $r(B) \geq 2\alpha \cdot \text{diam}(P)/\lambda$ be its radius. Since $(1 + \lambda/2\alpha)r(B) \geq r(B) + \text{diam}(P)$, the concentric ball of radius $(1 + \lambda/2\alpha)r(B)$ contains P . It follows that there is a cover for P that consists of a single ball and has cost at most

$$(1 + \lambda/2\alpha)^\alpha r(B)^\alpha \leq (1 + \lambda)\text{cost}(\text{opt}(P)) \leq (1 + \lambda)^l \text{cost}(\text{opt}(P)).$$

Case 2: There is no ball in $\text{opt}(P)$ whose radius is at least $2\alpha \cdot \text{diam}(P)/\lambda$. Let $\mathcal{Q}_0 \subseteq \text{opt}(P)$ contain those balls of radius at least $\text{diam}(P)/\gamma$. It follows from Lemma 3 that $|\mathcal{Q}_0| \leq c(\lambda, \gamma)$. Thus the algorithm considers a \mathcal{Q} with $\mathcal{Q} = \mathcal{Q}_0$. Fix this iteration. Also fix the partition $\{P_1, \dots, P_\tau\}$ of P computed by RAND-PARTITION(P). RAND-PARTITION ensures that $\text{diam}(P_i) \leq \text{diam}(P)/2$ for $1 \leq i \leq \tau$. Thus $\text{diam}(P'_i) \leq \text{diam}(P)/2$ and the level of each P'_i is at most $l - 1$. Hence by induction the expected value of $\text{cost}(\text{POINT-COVER}(P'_i))$ is at most $(1 + \lambda)^{l-1} \text{cost}(\text{opt}(P'_i))$.

Let $\mathcal{S}' = \text{opt}(P) \setminus \mathcal{Q}_0$. We argue below that the expected value of $\sum_{i=1}^\tau \text{cost}(\text{opt}(P'_i))$ is at most $(1 + \lambda)\text{cost}(\mathcal{S}')$. Assuming this, we have

$$\begin{aligned}
E[\text{cost}(\mathcal{Q}_0 \cup \bigcup_{i=1}^\tau \text{POINT-COVER}(P'_i))] &\leq \text{cost}(\mathcal{Q}_0) + (1 + \lambda)^{l-1} E[\sum_{i=1}^\tau \text{cost}(\text{opt}(P'_i))] \\
&\leq \text{cost}(\mathcal{Q}_0) + (1 + \lambda)^l \text{cost}(\mathcal{S}') \\
&\leq (1 + \lambda)^l \text{cost}(\text{opt}(P)).
\end{aligned}$$

XX:12 Approximate Clustering via Metric Partitioning

Thus $\text{POINT-COVER}(P)$ returns a solution whose expected cost is at most $(1+\lambda)^l \text{cost}(\text{opt}(P))$, as desired.

We now argue that the expected value of $\sum_{i=1}^{\tau} \text{cost}(\text{opt}(P'_i))$ is at most $(1+\lambda)\text{cost}(\mathcal{S}')$. Let \mathcal{B}_i consist of those balls in \mathcal{S}' that intersect P_i . For $B \in \mathcal{S}'$, let $\mu(B)$ denote the number of blocks in the partition $\{P_1, \dots, P_{\tau}\}$ that B intersects. Because \mathcal{B}_i is a cover for P'_i , we have $\text{cost}(\text{opt}(P'_i)) \leq \text{cost}(\mathcal{B}_i)$. Thus

$$\sum_{i=1}^{\tau} \text{cost}(\text{opt}(P'_i)) \leq \sum_{i=1}^{\tau} \text{cost}(\mathcal{B}_i) = \sum_{B \in \mathcal{S}'} \mu(B) \text{cost}(B).$$

By definition of \mathcal{Q}_0 , any ball $B \in \mathcal{S}' = \text{opt}(P) \setminus \mathcal{Q}_0$ has radius at most $\frac{\text{diam}(P)}{\gamma} = \frac{\lambda \cdot \text{diam}(P)}{c \log n}$, where c is the constant in Theorem 2. We may assume that $c \geq 16$ and hence $\frac{\lambda \cdot \text{diam}(P)}{c \log n} \leq \frac{\text{diam}(P)}{16 \log n}$. Theorem 2 now implies that

$$E[\mu(B)] \leq 1 + \frac{c \cdot r(B) \log n}{\text{diam}(P)} \leq 1 + \frac{c \log n}{\text{diam}(P)} \cdot \frac{\lambda \cdot \text{diam}(P)}{c \log n} = 1 + \lambda.$$

Thus the expected value of $\sum_{i=1}^{\tau} \text{cost}(\text{opt}(P'_i))$ is at most

$$\sum_{B \in \mathcal{S}'} E[\mu(B)] \text{cost}(B) \leq (1+\lambda) \sum_{B \in \mathcal{S}'} \text{cost}(B) = (1+\lambda) \text{cost}(\mathcal{S}'),$$

as claimed. \blacktriangleleft

We conclude that the expected cost of the cover returned by $\text{POINT-COVER}(X)$ is at most $(1+\lambda)^L \text{cost}(\text{opt}(X)) \leq (1+\varepsilon) \text{cost}(\text{opt}(X))$, since $\lambda = \varepsilon/2L$.

Now consider the time complexity of the algorithm. $\text{POINT-COVER}(P)$ makes $(mn)^{O(c(\lambda, \gamma))}$ direct recursive calls on subsets of diameter at most $\text{diam}(P)/2$. Thus the overall time complexity of $\text{POINT-COVER}(X)$ can be bounded by $(mn)^{O(c(\lambda, \gamma)L)}$. Plugging in $\lambda = \varepsilon/2L$, $\gamma = c \log n / \lambda$, and $c(\lambda, \gamma) = (9\alpha\gamma/\lambda)^{\alpha}$, we conclude

► **Theorem 5.** *There is an algorithm for MCC that runs in time $(mn)^{O(\frac{\alpha L^2 \log n}{\varepsilon^2})^{\alpha} L}$ and returns a cover whose expected cost is at most $(1+\varepsilon)$ times the optimal. Here L is 1 plus the logarithm of the aspect ratio of X , that is, the ratio of the maximum and minimum interpoint distances in the client set X .*

Using relatively standard techniques, which we omit here, we can pre-process the input to ensure that the ratio of the maximum and minimum interpoint distances in X is upper bounded by a polynomial in $\frac{mn}{\varepsilon}$. However, this affects the optimal solution by a factor of at most $(1+\varepsilon)$. After this pre-processing, we have $L = O(\log \frac{mn}{\varepsilon})$. Using the algorithm in Theorem 5 after the pre-processing, we obtain a $(1+\varepsilon)$ approximation with the quasi-polynomial running time $O(2^{\log^{O(1)} mn})$. Here the $O(1)$ hides a constant that depends on α and ε .

4 Algorithm for k -clustering

Recall that in k -clustering we are given a set X of points, a metric d on X , and a positive integer k . Let $|X| = n$. For $P \subseteq X$ and integer $\kappa \geq 0$, let $\text{opt}(P, \kappa)$ denote an optimal solution of κ -clustering for P (using balls whose center can be any point in X). We reuse the notions of $\text{level}(P)$, $\text{cost}(B)$ and $\text{cost}(\mathcal{B})$ from Section 3, for a point set P , a ball B , and a set \mathcal{B} of balls, respectively. Denote \mathcal{D} to be the set of balls such that each ball is centered at

a point of $y \in X$ and has radius $r = d(x, y)$ for some $x \in X$. We note that for any $P \subseteq X$, any ball in $\text{opt}(P, \kappa)$ must belong to this set. Note that $|\mathcal{D}| \leq n^2$.

To start with we prove a structure lemma for k -clustering.

► **Lemma 6.** *Let $P \subseteq X$, κ be a positive integer, and $\gamma \geq 1$. Then the number of balls in $\text{opt}(P, \kappa)$ of radius greater than or equal to $\text{diam}(P)/\gamma$ is at most $c(\gamma) := \gamma^\alpha$.*

Proof. Note that any ball centered at a point in P and having radius $\text{diam}(P)$ contains all the points of P . Now by definition of $\text{diam}(P)$ and \mathcal{D} , there is a point $x \in P$ such that the ball $B(x, \text{diam}(P)) \in \mathcal{D}$. Hence $\text{opt}(P, \kappa) \leq \text{diam}(P)^\alpha$.

Let t denote the number of balls in $\text{opt}(P, \kappa)$ of radius greater than or equal to $\text{diam}(P)/\gamma$. By optimality of $\text{opt}(P, \kappa)$, we have $t \cdot (\text{diam}(P)/\gamma)^\alpha \leq \text{diam}(P)^\alpha$. Thus $t \leq \gamma^\alpha$. ◀

Like in the case of MCC, we assume that the minimum distance between two points in X is 1. Let $L = 1 + \log(\text{diam}(X))$. We fix a parameter $\lambda = \varepsilon/6L$. Let $\gamma = \frac{c \log n}{\lambda}$, where c is the constant in Theorem 2.

We design a procedure $\text{CLUSTERING}(P, \kappa)$ (see Algorithm 3) that given a subset P of X and an integer κ , returns a set of at most $(1 + 3\lambda)^l \kappa$ balls whose union contains P , where $l = \text{level}(P)$. We overview this procedure, focussing on the differences from the procedure $\text{POINT-COVER}()$ used to solve the MCC problem. In $\text{CLUSTERING}(P, \kappa)$, $\text{RAND-PARTITION}(P)$ is called with $Z = X$ and $n = |X| \geq |P|$. We require two properties of the partition $\{P_1, \dots, P_\tau\}$ of P computed by $\text{RAND-PARTITION}(P)$. Let \mathcal{Q}_0 be the set containing the large balls of $\text{opt}(P, \kappa)$, that is, those with radius at least $\text{diam}(P)/\gamma$. Let $\mathcal{S}' = \text{opt}(P, \kappa) \setminus \mathcal{Q}_0$ denote the set of small balls, and let $\mathcal{S}'_i \subseteq \mathcal{S}'$ consist of those balls that contain at least one point in P_i that is not covered by \mathcal{Q}_0 . We would like (a) $\sum_{i=1}^\tau \text{cost}(\mathcal{S}'_i) \leq (1 + 3\lambda) \text{cost}(\mathcal{S}')$, and (b) $\sum_{i=1}^\tau |\mathcal{S}'_i| \leq (1 + 3\lambda) |\mathcal{S}'|$. Theorem 2 ensures that each of (a) and (b) holds in expectation. However, we would like both (a) and (b) to hold simultaneously, not just in expectation. For this reason, we try $\Theta(\log n)$ independent random partitions in Line 6, ensuring that with high probability, properties (a) and (b) hold for at least one of them.

Now let us fix one of these $\Theta(\log n)$ trials where we got a partition $\{P_1, \dots, P_\tau\}$ satisfying properties (a) and (b), and also fix an iteration in Line 9 where we have $\mathcal{Q} = \mathcal{Q}_0$. Let $P'_i \subseteq P_i$ be the points not covered by \mathcal{Q}_0 . For each $1 \leq i \leq \tau$ and $0 \leq \kappa_1 \leq (1 + 3\lambda)\kappa$, we set $\text{cluster}(P'_i, \kappa_1)$ to be the cover obtained by recursively invoking $\text{CLUSTERING}(P'_i, \kappa_1)$ (as in Line 13).

Let us call a tuple $(\kappa_1, \kappa_2, \dots, \kappa_\tau)$ of integers *valid* if $0 \leq \kappa_i \leq (1 + 3\lambda)(\kappa - |\mathcal{Q}_0|)$ and $\sum_{i=1}^\tau \kappa_i \leq (1 + 3\lambda)(\kappa - |\mathcal{Q}_0|)$. We would like to minimize $\sum_{i=1}^\tau \text{cost}(\text{cluster}(P'_i, \kappa_i))$ over all valid tuples $(\kappa_1, \kappa_2, \dots, \kappa_\tau)$. As there are too many valid tuples to allow explicit enumeration, we solve this optimization problem in Lines 14–21 via a dynamic programming approach.

This completes our overview. Our overall algorithm for k -clustering calls the procedure $\text{CLUSTERING}(X, k)$. Next we give the approximation bound on the cost of the solution returned by $\text{CLUSTERING}(P, \kappa)$.

► **Lemma 7.** *For any $P \subseteq X$ and an integer $\kappa \geq 1$, $\text{CLUSTERING}(P, \kappa)$ returns a solution consisting of at most $(1 + 3\lambda)^l \kappa$ balls and with probability at least $1 - \frac{|P|-1}{n^2}$, the cost of the solution is at most $(1 + 3\lambda)^l \text{cost}(\text{opt}(P, \kappa))$, where $l = \text{level}(P)$.*

Proof. We prove this lemma using induction on l . If $l = 0$, then $|P| \leq 1$ and $\text{CLUSTERING}(P, \kappa)$ returns an optimal solution, whose cost is $\text{cost}(\text{opt}(P, \kappa))$. Thus assume that $l \geq 1$ and the statement is true for subsets having level at most $l - 1$. If $|P|$ is smaller than the constant

Algorithm 3 CLUSTERING(P, κ)**Require:** A subset $P \subseteq X$, an integer κ .**Ensure:** A set of balls whose union contains the points in P .

```

1: if  $|P|$  is smaller than some constant  $\beta$  then
2:   return a minimum solution by checking all solutions with at most  $\min\{\kappa, \beta\}$  balls.
3:  $\text{sol} \leftarrow$  the best solution with one ball
4:  $\text{cost} \leftarrow \text{cost}(\text{sol})$ 
5:  $l \leftarrow \text{level}(P)$ 
6: for all  $2 \log_{3/2} n$  iterations do
7:   Let  $\{P_1, \dots, P_\tau\}$  be the set of nonempty subsets returned by RAND-PARTITION( $P$ )
8:   Let  $\mathcal{B}$  be the set of balls in  $\mathcal{D}$  having radius greater than  $\frac{\text{diam}(P)}{\gamma}$ 
9:   for each  $\mathcal{Q} \subseteq \mathcal{B}$  of size at most  $c(\gamma)$  do
10:    for  $i = 1$  to  $\tau$  do
11:      Let  $P'_i = \{p \in P_i \mid p \notin \bigcup_{B \in \mathcal{Q}} B\}$ 
12:      for each  $1 \leq i \leq \tau$  and  $0 \leq \kappa_1 \leq (1+3\lambda)\kappa$  do
13:         $\text{cluster}(P'_i, \kappa_1) \leftarrow \text{CLUSTERING}(P'_i, \kappa_1)$ 
14:      for  $i = 0$  to  $\tau - 1$  do
15:         $R_i \leftarrow \bigcup_{j=i+1}^\tau P'_j$ 
16:        for  $\kappa_1 = 0$  to  $(1+3\lambda)\kappa$  do
17:           $\text{cluster}(R_{\tau-1}, \kappa_1) \leftarrow \text{cluster}(P'_\tau, \kappa_1)$ 
18:          for all  $i = \tau - 2$  to  $0$  and  $0 \leq \kappa_1 \leq (1+3\lambda)\kappa$  do
19:             $\kappa'_{\min} \leftarrow \arg \min_{\kappa': 0 \leq \kappa' \leq \kappa_1} \text{cost}(\text{cluster}(P'_{i+1}, \kappa') \cup \text{cluster}(R_{i+1}, \kappa_1 - \kappa'))$ 
20:             $\text{cluster}(R_i, \kappa_1) \leftarrow \text{cluster}(P'_{i+1}, \kappa'_{\min}) \cup \text{cluster}(R_{i+1}, \kappa_1 - \kappa'_{\min})$ 
21:           $\mathcal{Q}' \leftarrow \mathcal{Q} \cup \text{cluster}(R_0, (1+3\lambda) \cdot (\kappa - |\mathcal{Q}|))$ 
22:          if  $|\mathcal{Q}'| \leq (1+3\lambda)^l \kappa$  and  $\text{cost}(\mathcal{Q}') < \text{cost}$  then
23:             $\text{cost} \leftarrow \text{cost}(\mathcal{Q}')$ 
24:             $\text{sol} \leftarrow \mathcal{Q}'$ 
25: return  $\text{sol}$ 

```

threshold β , CLUSTERING(P, κ) returns an optimal solution. So we may assume that $|P|$ is larger than this threshold.

Consider one of the $2 \log_{3/2} n$ iterations of CLUSTERING(P, κ). Fix the partition $\{P_1, \dots, P_\tau\}$ of P computed by RAND-PARTITION(P) in this iteration. Let \mathcal{Q}_0 be the set containing the balls of $\text{opt}(P, \kappa)$ with radius at least $\text{diam}(P)/\gamma$. It follows from Lemma 6 that $|\mathcal{Q}_0| \leq c(\gamma)$. Fix the choice $\mathcal{Q} = \mathcal{Q}_0$.

RAND-PARTITION(P) ensures that $\text{diam}(P_i) \leq \text{diam}(P)/2$ for $1 \leq i \leq \tau$. Thus $\text{diam}(P'_i) \leq \text{diam}(P)/2$ and the level of each P'_i is at most $l - 1$. Hence by induction, $\text{cluster}(P'_i, \kappa_1)$ contains at most $(1+3\lambda)^{l-1} \kappa_1$ balls and with probability at least $1 - \frac{|P'_i| - 1}{n^2}$, its cost is at most $(1+3\lambda)^{l-1} \text{cost}(\text{opt}(P'_i, \kappa_1))$ for $1 \leq \kappa_1 \leq (1+3\lambda)\kappa$.

Let $\mathcal{S}' = \text{opt}(P, \kappa) \setminus \mathcal{Q}_0$. Thus $|\mathcal{S}'| \leq \kappa - |\mathcal{Q}_0|$. We note that the union of the balls in \mathcal{S}' contains the points in $\bigcup_{i=1}^\tau P'_i$. Let $\mathcal{S}'_i \subseteq \mathcal{S}'$ be the set of balls that intersect with P'_i and thus the union of balls in \mathcal{S}'_i contains P'_i , where $1 \leq i \leq \tau$. We argue below that the value of $\sum_{i=1}^\tau |\mathcal{S}'_i|$ is at most $(1+3\lambda)|\mathcal{S}'|$ and the value of $\sum_{i=1}^\tau \text{cost}(\text{opt}(P'_i, |\mathcal{S}'_i|))$ is at most $(1+3\lambda)\text{cost}(\mathcal{S}')$ with probability at least $1/3$. Then the probability is at least $1 - \frac{1}{n^2}$ that corresponding to one of the $2 \log_{3/2} n$ iterations the value of $\sum_{i=1}^\tau |\mathcal{S}'_i|$ is at most $(1+3\lambda)|\mathcal{S}'|$ and the value of $\sum_{i=1}^\tau \text{cost}(\text{opt}(P'_i, |\mathcal{S}'_i|))$ is at most $(1+3\lambda)\text{cost}(\mathcal{S}')$. We refer to this event as the first good event. Now let us assume that the first good event occurs.

Note that by induction, $\text{cluster}(P'_i, |\mathcal{S}'_i|)$ contains at most $(1 + 3\lambda)^{l-1} |\mathcal{S}'_i|$ balls and with probability at least $1 - \frac{|P'_i| - 1}{n^2}$, its cost is at most $(1 + 3\lambda)^{l-1} \text{cost}(\text{opt}(P'_i, |\mathcal{S}'_i|))$ for $1 \leq i \leq \tau$. The probability that for every $1 \leq i \leq \tau$, the cost of $\text{cluster}(P'_i, |\mathcal{S}'_i|)$ is at most $(1 + 3\lambda)^{l-1} \text{cost}(\text{opt}(P'_i, |\mathcal{S}'_i|))$ and it contains at most $(1 + 3\lambda)^{l-1} |\mathcal{S}'_i|$ balls, is at least

$$\prod_{i=1}^{\tau} 1 - \frac{|P'_i| - 1}{n^2} \geq \prod_{i=1}^{\tau} 1 - \frac{|P_i| - 1}{n^2} \geq 1 - \frac{|P| - 2}{n^2}.$$

We refer to this event as the second good event. Assuming this second good event occurs, it follows by an easy induction on i that $\text{cluster}(R_i, \sum_{j=i+1}^{\tau} |\mathcal{S}'_j|)$ covers $R_i = \bigcup_{j=i+1}^{\tau} P'_j$ with at most $(1 + 3\lambda)^{l-1} \sum_{j=i+1}^{\tau} |\mathcal{S}'_j|$ balls and its cost is at most $(1 + 3\lambda)^{l-1} \sum_{j=i+1}^{\tau} \text{cost}(\text{opt}(P'_j, |\mathcal{S}'_j|))$. Thus $\text{cluster}(R_0, \sum_{i=1}^{\tau} |\mathcal{S}'_i|)$ covers $R_0 = \bigcup_{i=1}^{\tau} P'_i$ with at most $(1 + 3\lambda)^{l-1} \sum_{i=1}^{\tau} |\mathcal{S}'_i|$ balls and its cost is at most $(1 + 3\lambda)^{l-1} \sum_{i=1}^{\tau} \text{cost}(\text{opt}(P'_i, |\mathcal{S}'_i|))$. Hence $\mathcal{Q}_0 \cup \text{cluster}(R_0, \sum_{i=1}^{\tau} |\mathcal{S}'_i|)$ covers P . Now assuming both the good events occur, we have

$$\begin{aligned} |\mathcal{Q}_0 \cup \text{cluster}(R_0, \sum_{i=1}^{\tau} |\mathcal{S}'_i|)| &\leq |\mathcal{Q}_0| + \sum_{i=1}^{\tau} (1 + 3\lambda)^{l-1} |\mathcal{S}'_i| \\ &\leq |\mathcal{Q}_0| + (1 + 3\lambda)^{l-1} \sum_{i=1}^{\tau} |\mathcal{S}'_i| \\ &\leq |\mathcal{Q}_0| + (1 + 3\lambda)^l |\mathcal{S}'| \\ &\leq |\mathcal{Q}_0| + (1 + 3\lambda)^l (\kappa - |\mathcal{Q}_0|) \\ &\leq (1 + 3\lambda)^l \kappa. \end{aligned}$$

$$\begin{aligned} \text{cost}(\mathcal{Q}_0 \cup \text{cluster}(R_0, \sum_{i=1}^{\tau} |\mathcal{S}'_i|)) &\leq \text{cost}(\mathcal{Q}_0) + (1 + 3\lambda)^{l-1} \sum_{i=1}^{\tau} \text{cost}(\text{opt}(P'_i, |\mathcal{S}'_i|)) \\ &\leq \text{cost}(\mathcal{Q}_0) + (1 + 3\lambda)^l \text{cost}(\mathcal{S}') \\ &\leq (1 + 3\lambda)^l \text{cost}(\text{opt}(P, \kappa)). \end{aligned}$$

The probability that both the good events occur is at least

$$(1 - \frac{|P| - 2}{n^2}) \cdot (1 - \frac{1}{n^2}) \geq 1 - \frac{|P| - 1}{n^2}.$$

Hence the statement of the lemma follows by noting that the value of $\sum_{i=1}^{\tau} |\mathcal{S}'_i|$ in $\text{cluster}(R_0, \sum_{i=1}^{\tau} |\mathcal{S}'_i|)$ is at most $(1 + 3\lambda) \cdot (\kappa - |\mathcal{Q}_0|)$.

We now argue that the probability that $\sum_{i=1}^{\tau} |\mathcal{S}'_i|$ is greater than $(1 + 3\lambda) |\mathcal{S}'|$ is at most $1/3$ and the probability that $\sum_{i=1}^{\tau} \text{cost}(\text{opt}(P'_i, |\mathcal{S}'_i|))$ is greater than $(1 + 3\lambda) \text{cost}(\mathcal{S}')$ is at most $1/3$. Then using union bound, the probability that $\sum_{i=1}^{\tau} |\mathcal{S}'_i|$ is at most $(1 + 3\lambda) |\mathcal{S}'|$ and $\sum_{i=1}^{\tau} \text{cost}(\text{opt}(P'_i, |\mathcal{S}'_i|))$ is at most $(1 + 3\lambda) \text{cost}(\mathcal{S}')$ is at least $1/3$.

For $B \in \mathcal{S}'$, let $\mu(B)$ denote the number of blocks in the partition $\{P_1, \dots, P_{\tau}\}$ that non-terminally intersect B . We note that at most one block can terminally intersect B . Because \mathcal{S}'_i is a cover for P'_i with $|\mathcal{S}'_i|$ balls, we have $\text{cost}(\text{opt}(P'_i, |\mathcal{S}'_i|)) \leq \text{cost}(\mathcal{S}'_i)$. Thus

$$\sum_{i=1}^{\tau} \text{cost}(\text{opt}(P'_i, |\mathcal{S}'_i|)) \leq \sum_{i=1}^{\tau} \text{cost}(\mathcal{S}'_i) = \sum_{B \in \mathcal{S}'} (1 + \mu(B)) \text{cost}(B) = \text{cost}(\mathcal{S}') + \sum_{B \in \mathcal{S}'} \mu(B) \text{cost}(B).$$

and,

$$\sum_{i=1}^{\tau} |\mathcal{S}'_i| \leq \sum_{B \in \mathcal{S}'} (1 + \mu(B)) = |\mathcal{S}'| + \sum_{B \in \mathcal{S}'} \mu(B).$$

XX:16 Approximate Clustering via Metric Partitioning

By definition of \mathcal{Q}_0 , any ball $B \in \mathcal{S}' = \text{opt}(P, \kappa) \setminus \mathcal{Q}_0$ has radius at most $\frac{\text{diam}(P)}{\gamma} = \frac{\lambda \cdot \text{diam}(P)}{c \log n}$, where c is the constant in Theorem 2. We may assume that $c \geq 16$ and hence $\frac{\lambda \cdot \text{diam}(P)}{c \log n} \leq \frac{\text{diam}(P)}{16 \log n}$. Theorem 2 now implies that

$$E[\mu(B)] \leq \frac{c \cdot r(B) \log n}{\text{diam}(P)} \leq \frac{c \log n}{\text{diam}(P)} \cdot \frac{\lambda \cdot \text{diam}(P)}{c \log n} = \lambda.$$

Using linearity of expectation,

$$E\left[\sum_{B \in \mathcal{S}'} \mu(B)\right] = \sum_{B \in \mathcal{S}'} E[\mu(B)] \leq \lambda \cdot |\mathcal{S}'|.$$

Now by Markov's inequality,

$$\Pr\left[\sum_{B \in \mathcal{S}'} \mu(B) > 3\lambda \cdot |\mathcal{S}'|\right] \leq 1/3.$$

It follows that,

$$\begin{aligned} & \Pr\left[\sum_{i=1}^{\tau} |\mathcal{S}'_i| > (1 + 3\lambda)|\mathcal{S}'|\right] \\ & \leq \Pr\left[|\mathcal{S}'| + \sum_{B \in \mathcal{S}'} \mu(B) > (1 + 3\lambda)|\mathcal{S}'|\right] \\ & \leq \Pr\left[\sum_{B \in \mathcal{S}'} \mu(B) > 3\lambda \cdot |\mathcal{S}'|\right] \leq 1/3. \end{aligned}$$

Again using linearity of expectation,

$$E\left[\sum_{B \in \mathcal{S}'} \mu(B) \text{cost}(B)\right] = \sum_{B \in \mathcal{S}'} E[\mu(B)] \text{cost}(B) \leq \lambda \cdot \text{cost}(\mathcal{S}').$$

Now by Markov's inequality,

$$\Pr\left[\sum_{B \in \mathcal{S}'} \mu(B) \text{cost}(B) > 3\lambda \cdot \text{cost}(\mathcal{S}')\right] \leq 1/3.$$

It follows that,

$$\begin{aligned} & \Pr\left[\sum_{i=1}^{\tau} \text{cost}(\text{opt}(P'_i, |\mathcal{S}'_i|)) > (1 + 3\lambda) \text{cost}(\mathcal{S}')\right] \\ & \leq \Pr\left[\text{cost}(\mathcal{S}') + \sum_{B \in \mathcal{S}'} \mu(B) \text{cost}(B) > (1 + 3\lambda) \text{cost}(\mathcal{S}')\right] \\ & \leq \Pr\left[\sum_{B \in \mathcal{S}'} \mu(B) \text{cost}(B) > 3\lambda \cdot \text{cost}(\mathcal{S}')\right] \leq 1/3. \end{aligned}$$

◀

Since $\lambda = \varepsilon/6L$, $(1 + 3\lambda)^L \leq 1 + \varepsilon$. Thus we conclude that with probability at least $1 - \frac{1}{n}$, $\text{CLUSTERING}(X, k)$ returns a solution with at most $(1 + \varepsilon)k$ balls whose cost is at most $(1 + \varepsilon) \text{cost}(\text{opt}(X, k))$.

Now consider the time complexity of the algorithm. $\text{CLUSTERING}(P, \kappa)$ makes $n^{O(c(\gamma))}$ direct recursive calls on subsets of diameter at most $\text{diam}(P)/2$. Thus the overall time complexity of $\text{CLUSTERING}(X, k)$ can be bounded by $n^{O(c(\gamma)L)}$. Plugging in $\lambda = \varepsilon/6L$, $\gamma = c \log n / \lambda$, and $c(\gamma) = \gamma^\alpha$, we conclude

► **Theorem 8.** *There is a randomized algorithm for k -clustering that runs in time $n^{O((\frac{L \log n}{\epsilon})^\alpha L)}$ and with probability at least $1 - \frac{1}{n}$ returns a solution with at most $(1 + \epsilon)k$ balls whose cost is at most $(1 + \epsilon)$ times the optimal. Here L is 1 plus the logarithm of the aspect ratio of X , that is, the ratio of the maximum and minimum interpoint distances in the set X .*

5 Inapproximability Result

In this section we present an inapproximability result which complements the result in Section 3. In particular here we consider the case when α is not a constant. The heart of this result is a reduction from the dominating set problem. Given a graph $G = (V, E)$, a dominating set for G is a subset V' of V such that for any vertex $v \in V \setminus V'$, v is connected to at least one vertex of V' by an edge in E . The dominating set problem is defined as follows.

Dominating Set Problem (DSP)

INSTANCE: Graph $G = (V, E)$, positive integer $k \leq |V|$.

QUESTION: Is there a dominating set for G of size at most k ?

The following inapproximability result is proved by Kann [20].

► **Theorem 9.** *There is a constant $c > 0$ such that there is no polynomial-time $c \log |V|$ -factor approximation algorithm for DSP assuming $\mathcal{P} \neq \mathcal{NP}$.*

The following theorem shows an inapproximability bound for MCC when $\alpha \geq \log |X|$.

► **Theorem 10.** *For $\alpha \geq \log |X|$, no polynomial time algorithm for MCC can achieve an approximation factor better than $c \log |X|$ assuming $\mathcal{P} \neq \mathcal{NP}$.*

Proof. To prove this theorem we show a reduction from DSP. Given an instance $(G = (V, E), k)$ of DSP we construct an instance of MCC. The instance of MCC consists of two sets of points X (clients) and Y (servers), and a metric d defined on $X \cup Y$. Let $V = \{v_1, v_2, \dots, v_n\}$, where $n = |V|$. For each $v_i \in V$, Y contains a point y_i and X contains a point x_i . For any point $p \in X \cup Y$, $d(p, p) = 0$. For $i, j \in [n]$, $d(x_i, y_j)$ is 1 if $i = j$ or the edge $(v_i, v_j) \in E$, and $d(x_i, y_j)$ is 3 otherwise. For $i, j \in [n]$ such that $i \neq j$, we set $d(x_i, x_j) = d(y_i, y_j) = 2$.

Consider two nonadjacent vertices v_i and v_j . For any $x_t \in X$ such that $t \neq i, j$, $d(x_i, x_t) + d(x_t, y_j) \geq 3$. Similarly, for any $y_t \in Y$ such that $t \neq i, j$, $d(x_i, y_t) + d(y_t, y_j) \geq 3$. Thus d defines a metric. Next we will prove that G has a dominating set of size at most k iff the cost of covering the points in X using the balls around the points in Y is at most k .

Suppose G has a dominating set J of size at most k . For each vertex $v_j \in J$, build a radius 1 ball around y_j . We return this set of balls \mathcal{B} as the solution of MCC. Now consider any point $x_i \in X$. If $v_i \in J$, then x_i is covered by the ball around y_i . Otherwise, there must be a vertex $v_j \in J$ such that $(v_i, v_j) \in E$. Then $d(x_i, y_j)$ is 1 and x_i is covered by the ball around y_j . Hence \mathcal{B} is a valid solution of MCC with cost at most k .

Now suppose there is a solution \mathcal{B} of MCC with cost at most k . If $k > |X|$, then V is a dominating set for G of size $|X| < k$. If $k \leq |X|$, our claim is that the radius of each ball in \mathcal{B} is 1. Suppose one of the balls B has a radius more than 1. Then the way the instance of MCC is created the radius should be at least 3. Hence $k \geq 3^\alpha \geq 3^{\log |X|} > |X|$, which is a contradiction. Now consider the set of vertices J corresponding to the centers of balls in \mathcal{B} . It is not hard to see that J is a dominating set for G of size at most k .

Let OPT be the cost of any optimal solution of MCC for the instance (X, Y, d) . Then by the properties of this reduction the size of any minimum dominating set for G is OPT. Thus if

there is an approximation algorithm for MCC that gives a solution with cost $(c \log |X|) \cdot \text{OPT}$, then using the reduction we can produce a dominating set of size $(c \log |V|) \cdot \text{OPT}$. Then from Theorem 9 it follows that $\mathcal{P} = \mathcal{NP}$. This completes the proof of our theorem. \blacktriangleleft

6 Conclusions

One generalization of the MCC problem that has been studied [9, 8] includes fixed costs for opening the servers. As input, we are given two point sets X (clients) and Y (servers), a metric on $Z = X \cup Y$, and a *facility cost* $f_y \geq 0$ for each server $y \in Y$. The goal is to find a subset $Y' \subseteq Y$, and a set of balls $\{B_y \mid y \in Y' \text{ and } B_y \text{ is centered at } y\}$ that covers X , so as to minimize $\sum_{y \in Y'} (f_y + r(B_y)^\alpha)$. It is not hard to see that our approach generalizes in a straightforward way to give a $(1 + \varepsilon)$ approximation to this problem using quasi-polynomial running time. To keep the exposition clear, we have focussed on the MCC rather than this generalization.

The main open problem that emerges from our work is whether there one can obtain a $(1 + \varepsilon)$ -approximation for the k -clustering problem in quasi-polynomial time.

Acknowledgements. We would like to thank an anonymous reviewer of an earlier version of this paper for suggestions that improved the guarantees and simplified the proof of Theorem 2. We also thank other reviewers for their feedback and pointers to the literature.

References

- 1 Ittai Abraham, Yair Bartal, and Ofer Neimany. Advances in metric embedding theory. In *Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*, pages 271–286. ACM, 2006.
- 2 Helmut Alt, Esther M. Arkin, Hervé Brönnimann, Jeff Erickson, Sándor P. Fekete, Christian Knauer, Jonathan Lenchner, Joseph S. B. Mitchell, and Kim Whittlesey. Minimum-cost coverage of point sets by disks. In *Proceedings of the 22nd ACM Symposium on Computational Geometry, Sedona, Arizona, USA, June 5-7, 2006*, pages 449–458, 2006.
- 3 Sanjeev Arora, Prabhakar Raghavan, and Satish Rao. Approximation schemes for euclidean k -medians and related problems. In *Proceedings of the Thirtieth Annual ACM Symposium on Theory of Computing*, STOC '98, pages 106–113, New York, NY, USA, 1998. ACM.
- 4 Sayan Bandyapadhyay and Kasturi R. Varadarajan. On variants of k -means clustering. In *32nd International Symposium on Computational Geometry, SoCG 2016, June 14-18, 2016, Boston, MA, USA*, pages 14:1–14:15, 2016.
- 5 Yair Bartal. Probabilistic approximations of metric spaces and its algorithmic applications. In *37th Annual Symposium on Foundations of Computer Science, FOCS '96, Burlington, Vermont, USA, 14-16 October, 1996*, pages 184–193. IEEE Computer Society, 1996.
- 6 Yair Bartal. Graph decomposition lemmas and their role in metric embedding methods. In *Algorithms-ESA 2004*, pages 89–97. Springer, 2004.
- 7 Babak Behsaz and Mohammad R. Salavatipour. On minimum sum of radii and diameters clustering. In *Algorithm Theory - SWAT 2012 - 13th Scandinavian Symposium and Workshops, Helsinki, Finland, July 4-6, 2012. Proceedings*, pages 71–82, 2012.
- 8 Vittorio Bilò, Ioannis Caragiannis, Christos Kaklamanis, and Panagiotis Kanellopoulos. Geometric clustering to minimize the sum of cluster sizes. In *Algorithms - ESA 2005, 13th*

- Annual European Symposium, Palma de Mallorca, Spain, October 3-6, 2005, Proceedings*, pages 460–471, 2005.
- 9 Moses Charikar and Rina Panigrahy. Clustering to minimize the sum of cluster diameters. *J. Comput. Syst. Sci.*, 68(2):417–441, 2004.
 - 10 Vincent Cohen-Addad, Philip N. Klein, and Claire Mathieu. Local search yields approximation schemes for k-means and k-median in euclidean and minor-free metrics. In *FOCS*, to appear, 2016.
 - 11 Srinivas Doddi, Madhav V. Marathe, S. S. Ravi, David Scot Taylor, and Peter Widmayer. Approximation algorithms for clustering to minimize the sum of diameters. *Nord. J. Comput.*, 7(3):185–203, 2000.
 - 12 Thomas Erlebach, Klaus Jansen, and Eike Seidel. Polynomial-time approximation schemes for geometric intersection graphs. *SIAM J. Comput.*, 34(6):1302–1323, 2005.
 - 13 Jittat Fakcharoenphol, Satish Rao, and Kunal Talwar. A tight bound on approximating arbitrary metrics by tree metrics. *J. Comput. Syst. Sci.*, 69(3):485–497, 2004.
 - 14 Ari Freund and Dror Rawitz. Combinatorial interpretations of dual fitting and primal fitting. In *Approximation and Online Algorithms, First International Workshop, WAOA 2003, Budapest, Hungary, September 16-18, 2003, Revised Papers*, pages 137–150, 2003.
 - 15 Matt Gibson, Gaurav Kanade, Erik Krohn, Imran A. Pirwani, and Kasturi R. Varadarajan. On metric clustering to minimize the sum of radii. *Algorithmica*, 57(3):484–498, 2010.
 - 16 Matt Gibson, Gaurav Kanade, Erik Krohn, Imran A. Pirwani, and Kasturi R. Varadarajan. On clustering to minimize the sum of radii. *SIAM J. Comput.*, 41(1):47–60, 2012.
 - 17 Sarel Har-Peled. Geometric approximation algorithms, 2011.
 - 18 Dorit S. Hochbaum and Wolfgang Maass. Approximation schemes for covering and packing problems in image processing and VLSI. *J. ACM*, 32(1):130–136, 1985.
 - 19 Lior Kamra, Robert Krauthgamer, and Huy L. Nguyen. Cutting corners cheaply, or how to remove steiner points. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1029–1040. SIAM, 2014.
 - 20 Viggo Kann. On the approximability of np-complete optimization problems. *PhD thesis, Department of Numerical Analysis and Computing Science, Royal Institute of Technology, Stockholm*.
 - 21 Nissan Lev-Tov and David Peleg. Polynomial time approximation schemes for base station coverage with minimum total radii. *Computer Networks*, 47(4):489–501, 2005.
 - 22 Satu Elisa Schaeffer. Graph clustering. *Computer Science Review*, 1(1):27–64, 2007.
 - 23 Mohammad Salavatipour, Zachary Friggstad, and Mohsen Rezapour. Local search yields a ptas for k -means in doubling metrics. In *FOCS*, to appear, 2016.